AntMonitor: System and Applications

Anastasia Shuba, Anh Le, Emmanouil Alimpertis, Minas Gjoka, Athina Markopoulou
{ashuba, anh.le, ealimper, mgjoka, athina}@uci.edu
University of California, Irvine

arXiv:1611.04268v1 [cs.NI] 14 Nov 2016

ABSTRACT

In this paper, we present a complete system for on-device passive monitoring, collection, and analysis of fine-grained, large-scale packet measurements from mobile devices. First, we describe the design and implementation of AntMonitor as a user-space mobile app based on a VPN-service but only on the device (without the need to route through a remote VPN server) and using only the minimum resources required. We evaluate our prototype and show that it significantly outperforms prior state-of-the-art approaches: it achieves throughput of over 90 Mbps downlink and 65 Mbps uplink, which is 2x and 8x faster than mobile-only baselines and is 94% of the throughput without VPN, while using 2–12x less energy. Second, we show that AntMonitor is uniquely positioned to serve as a platform for passive on-device mobile network monitoring and to enable a number of applications, including: (i) real-time detection and prevention of private information leakage from the device to the network; (ii) passive network performance monitoring; and (iii) application classification and user profiling. We showcase preliminary results from a pilot user study at a university campus.

1. INTRODUCTION

In addition to the large and ever-increasing volume of mobile network traffic [1], mobile devices are used today for a range of personal activities (from communication to financial transactions) and have access to personally identifiable information (PII). User behavior and third-party activity eventually manifest themselves as packets transmitted over the mobile network. Therefore, looking at the network activity on a mobile device provides a unique vantage point for monitoring and inferring a wealth of information, including PII leaks, network performance and user behavior.

There is a rich body of literature [2, 3, 4, 5, 6, 7, 8, 9] which studies network traffic traces, and the approaches typically fall into one of two categories: either large-scale but coarse-grained traces obtained in the middle of the network, *i.e.*, traces from Internet Service Providers (ISP) [2, 3], or fine-grained but small-scale traces from a limited set of users [4, 5]. These limitations, privacy and security concerns, and performance bottlenecks have hindered progress.

In this paper, we present AntMonitor: a system for on-device passive collection and analysis of fine-grained, large-scale, mobile network measurements. We argue that AntMonitor is well-positioned to become a high-performance passive monitoring tool for crowdsourcing a range of mobile network measurements, as it combines the following desired properties: (i) it is easy to install (it does not require administrative privileges) and use (it runs as a service app in the background); (ii) it scales well with the number of users thanks to its on-device design; (iii) it provides users with fine-grained control of which data to monitor or log; (iv) it supports real-time analysis on the device and/or off-line analysis on a log server, but does not intercept packets in middleboxes and (iv) it allows for semantic-rich annotation of packet traces with contextual information from the device.

We start by presenting the design and evaluation of AntMonitor. AntMonitor is a user-space mobile app, based on a VPN service (which intercepts all packets without rooting the phones) and routes packets directly to their destination (by translating all connections on the device, without redirecting them through a remote VPN server).[1] Our software architecture uses the minimum number of resources to achieve the highest performance, including (i) the minimum number of threads for network routing (two: one for reading/writing from/to the TUN) and (ii) the minimum number of sockets (by multiplexing UDP sockets and carefully managing TCP sockets), so as to not interfere with the performance of the foreground apps. Thanks to these and other design choices, AntMonitor significantly outperforms all prior state-of-the-art approaches, in terms of throughput and energy, without significantly impacting CPU and memory usage. More specifically, experiments show that AntMonitor achieves 2x (downlink) and 8x (uplink) throughput of state-of-the-art mobile-only approaches, namely Privacy Guard [10] and Haystack [11]. The achieved downlink throughput is also at 94% of the throughput without VPN, and almost double the throughput of comparable client-server approaches; all while using 2–12x less energy. We believe that this performance advantage is crucial for the success-

[1] Throughout the paper, we will refer to this approach as *mobile-only*, and to using a remote VPN server as *client-server*. In both cases, a separate logging server may be used for uploading logs from the device for subsequent offline analysis.

ful adoption of AntMonitor by users. However much users may care about added-value services, they are unlikely to install apps that slow down their phones or drain their battery.

Then, we demonstrate that AntMonitor naturally lends itself as a platform for a range of applications that build on top of passive on-device network monitoring, which can be of interest to individual users, network operators, and researchers. First, we use AntMonitor to detect and prevent in real-time *leakage of private information (PIIs)* from the device to the network. We show that AntMonitor was able to detect a large volume of PII leaks and the destinations where this information goes to, and we provide visualization of this information on the mobile device (in real-time) and/or on our LogServer (off-line) [12]. Second, we use AntMonitor for *passive performance measurements* network-wide (*e.g.*, network performance maps) as well as per-user (usage profiles) or per-app. This information comes at no additional overhead (*e.g.*, no active probing is needed) since AntMonitor touches every packet transmitted over the network, and can provide insights into network usage and provisioning as well as fine-grained information. Third, we use the packet traces collected by AntMonitor, extract features only from TCP/IP headers, annotate them with rich contextual information available on the device (such as location, time, foreground/background apps, etc.) to train *machine learning models*. We use these models for traffic classification of flows to applications (achieving higher classification accuracy than state-of-art methods that use HTTP payloads [13]) as well as for user profiling based on minimal information. We present results from a pilot user study at a university campus to demonstrate the capabilities of AntMonitor and its enabling potential for these applications.

The structure of the rest of the paper is as follows. Section 2 presents related work. Section 3 describes the objectives, design, and implementation of the AntMonitor system. Section 4 presents performance evaluation and comparison to state-of-the-art approaches. Section 5 describes the applications of AntMonitor to three domains, namely: privacy leak detection and prevention (Section 5.1); passive monitoring of network performance (Section 5.2); and application and user classification (Section 5.3). Section 6 concludes and discusses directions for future work.

2. RELATED WORK

2.1 VPN-based Mobile Network Monitoring

The approach we follow in AntMonitor is passive monitoring on the device, guided by the design objectives outlined in Section 3.1. In particular, the only way to intercept every packet in and out of the mobile device, while running in user space (without root or custom OS), today, is to establish a VPN service on the device. There are two VPN approaches: client-server and mobile-only, described below.

In **Client-Server VPN** approaches, packets are tunneled from the VPN client on the mobile device to a remote VPN server, where they can be processed or logged. A representative of this approach is Meddle [13], which builds on top of the StrongSwan [14] VPN software. Disadvantages of this approach include the fact that packets are routed through a middle server thus posing additional delay and privacy concerns, lack of client-side annotation (thus no ground truth available at the server), and potentially complex control mechanisms (the client has to communicate the selections of functionalities, *e.g.*, ad blocking, to the server). In [15], a client-server system was proposed that remedied the latter two disadvantages by building a custom client app. An advantage of the client-server approach is that it can be combined with other VPN and proxy services (*e.g.*, encryption, private browsing), and be offered by ISPs as an added-value service.

In **Mobile-Only VPN** approaches, the client establishes a VPN service on the phone to intercept all IP packets and does not require a VPN server for routing. It extracts the content of captured outgoing packets and sends them through newly created protected UDP/TCP sockets [16] to reach Internet hosts; and vice versa for incoming packets. This approach may have high overhead due to this layer-3 to layer-4 translation, the need to maintain state per connection and additional processing per packet. If not carefully implemented, this approach can significantly affect network throughput: for example, see the poor performance of tPacketCapture [17] – an application currently available on Google Play that utilizes this mobile-only approach. Two state-of-the-art representatives of the mobile-only approach are Haystack [11, 18] and Privacy Guard [10]. They both focus on applying and optimizing their systems for detection of PII leaks. Our evaluation in Section 4 shows that AntMonitor can achieve 2x and 8x their downlink and uplink throughput.

2.2 Other Monitoring Approaches

Work on monitoring network traffic generated by mobile devices can be roughly classified according to the vantage point and measurement approach.

Rooted phones and OS approaches. Using a custom OS or a rooted phone one can get access to fine-grained information on the device, including network traffic. Phonelab [7] and others [5, 4, 9] use packet capture APIs such as tcpdump or iptables-log. These are powerful but inherently limited to small scale-deployment as the overwhelming majority of users do not have rooted phones, and wireless providers and phone manufacturers strongly discourage rooting. The same limitation applies to OS-based approaches, including: TaintDroid that uses a custom OS, [19], MockDroid [20] and AppFence [21] that dynamically intercept any permission request to certain resources, AndroidLeaks [22] and PiOS [23] that use static analysis.

Active Measurements from Mobile Devices. There are mobile apps, developed by researchers Netalyzr [24], Mobilyzer [25] or the industry (*e.g.*, Speedtest, CarrierIQ or Tutella), to perform active network measurements of various metrics (throughput, latency, RSS) from the mobile de-

vice. They run in user space, without root, and allow for accurate measurements. However, care must be given to avoid burdening the device's resources and crowdsourcing is often used to distribute the load (see Section 2.3).

Passive Monitoring Inside the Network. ISPs and other organizations sometimes can capture packets on links in the middle of their networks, *e.g.* at an ISP's or other organization's network [26, 27, 28]. Researchers typically analyze network traces collected by others (*e.g.* large tier-1 networks [2] or from university campus WiFi networks [3]). Limitations of this approach include that (i) it only captures traffic going through the particular measurement point and (ii) it has access only to packet headers (payload is increasingly encrypted), not to ground truth or semantic-rich information (*e.g.* apps that produced the packets).

2.3 Applications of Network Monitoring

Privacy Leaks Detection. Next, we review work on detecting private data leaking out of a device, which is related to our first application in Section 5.1. Some approaches require a custom OS or rooting the phone, such as as TaintDroid [19] was one of the early tools built to identify privacy leaks in real-time, and others reviewed in Section 2.2 .

Another approach is to allow the user to define strings (*e.g.*, IMEI, device id, email, or any string corresponding to sensitive information that the user wants to protect) and then monitor for potential leaking of that information from the device to the network. AntMonitor (as well as Haystack and Privacy Guard) follow this approach: they monitor, on the device itself, the payload of all outgoing packets, searching for the predefined strings. In order to detect leaked strings in encrypted traffic, all three tools need a TLS proxy to first decrypt the traffic before string matching. Although the goal is the same, implementation matters: AntMonitor is currently the only tool that can prevent (*i.e.*, block or hash the private string), in addition to detection, on the mobile device without root privileges; AntMonitor and Privacy Guard can perform real-time detection, while Haystack does not.

Recon [29] also inspects packets for privacy leakage, but because it builds on top of Meddle [13], all packet processing (including privacy leaks detection) happens not on the device itself but on the Meddle server, with all the advantages and disadvantages of a client-server VPN discussed in Section 2.1. Recon is also the first to use machine learning to identify flows that leak private data without prior knowledge of the users' PII, based on HTTP features and training on user feedback, as well as on ground truth, manually obtained. This approach is also applicable to AntMonitor.

Performance Monitoring. Typically, network performance measurements (*e.g.* signal strength, throughput, delay) are outsourced to companies (some of them drive around and measure the networks) and/or are crowdsourced from individual users. Examples of *third-party companies* include: Carrier IQ [30], which is embedded in the firmware of over 150M smartphones and reports network information, signal strength and the users' location; Tutella, which provides a network performance SDK that can be embedded in other mobile applications; and Mobilyzer [25], which provides an open platform for controllable mobile network measurements. Examples of *crowdsourcing* projects include Speedtest [31], OpenSig-nals [32], Sensorly [33], and RootMetrics [34]. These companies make mobile applications that allow users to perform and report active measurements; the companies use that data to release performance reports [35] and awards [36] for cellular and Wi-Fi at points of interest (*e.g.*, metro areas, airports, sports venues, etc). Work in [8] analyzed crowdsourced data from Speedtest in order to compare the performance of cellular and Wi-Fi networks in large metro areas. Work in [37] released an app for crowdsourcing measurements of throughput and latency over LTE and Wi-Fi.

Similar to some of these approaches, AntMonitor crowdsources measurements from an end-user app. Contrary to the above apps, it can passively infer network performance metrics (e.g. throughput), without sending active probes. It can be used to create both network-wide performance maps and user-specific statistics, as discussed in Section 5.2.

Learning. Several papers [38, 39, 2, 40] perform *app classification* of flows by building signatures from unencrypted HTTP sessions in network traces: in [2], the HTTP User-Agent field was used to map flows into apps; in [38], HTTP header key-value pairs were used to build unique app signatures that operate on a per-flow basis; in [39], more flows could be identified by expanding the usage of tokens in the HTTP headers (beyond HTTP request data) and by propagating the identification of a flow mapped to a specific app to other flows that occur at the same time; in [40], conjunctive rules are constructed from HTTP header data. All these methods rely on HTTP headers whereas we perform app classification using only TCP/IP headers in Section 5.3. Early work on behavioral analysis classified protocols based on packet headers: graphlets [41], profiling tend-hosts [42], traffic dispersion graphs [43], subflows [44]. Prior work on *user profiling* includes: [45], which uses HMM classifiers on NetFlow records to fingerprint users behind NAT, and [46], which re-identifies users over different web sessions.

3. THE ANTMONITOR SYSTEM

3.1 Design Rationale

Here we describe the main objectives of AntMonitor and the key design choices made to meet the objectives.

Objective 1: Large-Scale Measurements: AntMonitor is intended for crowdsourcing data from a large number of users, which poses a number of system requirements. First, the app on the mobile device must run without administrative privileges (root access). To that end, we use the public Virtual Private Network (VPN) API [16] provided by the Android OS (version 4.0+), which runs on more than 95% of Android devices [47]. Second, in order for a large number of mobile users to adopt it, user experience must not be af-

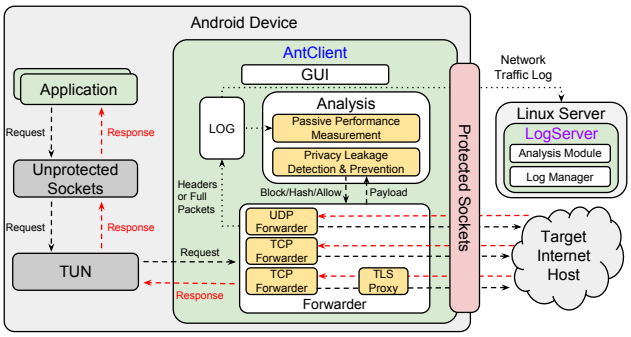

Figure 1: The AntMonitor Architecture. It consists of the AntMonitor App (on the device) and a remote LogServer.

fected: monitoring must occur seamlessly in the background while the user continues to use the mobile device as usual, and the overhead on the device must be negligible in terms of network throughput, CPU, battery, and data cost. Third, the performance must scale with the number of users, which makes a strong case for a mobile-only design.

Objective 2: Making it Attractive for Users: There must also be incentives for users to participate and AntMonitor is designed with the capability to offer users a variety of services. The current prototype offers enhanced privacy protection (*e.g.*, preventing leakage of private information) and visualizations that help users understand where their data flows (*e.g.* see Fig. 2(f)). Additional services can be implemented completely on the mobile, such as enhanced wireless network performance by switching among available networks; see Sec. 5.2. The AntMonitor App is also designed to provide users with control over which data they choose to monitor and log, including which applications to monitor, and whether to log full packets or headers only.

Objective 3: Fine-Grained Information: AntMonitor supports full-packet capture of incoming and outgoing traffic in PCAP Next Generation format [48], which allows to append arbitrary information alongside the raw packets. This is important because, in many cases, this contextual information may only be collected accurately on the mobile at the time of packet capture, and can play a critical role in subsequent analyses. Examples of such contextual information include names of apps that generate the packets (thus providing ground truth for application classification, see Section 5.3), location, background apps, and information about the network used (*e.g.* WiFi or LTE, network speed, *etc.*).

3.2 System Design

To support the above design objectives, AntMonitor is designed to provide the following four main functionalities: traffic interception, routing, logging, and analysis. The rationale for the first two is described here and implementation and optimizations for all four are in Sec. 3.3 and 3.5.

Traffic Interception. The AntMonitor App establishes a VPN service on the device that runs seamlessly in the background. The service is able to intercept all outgoing and incoming IP datagrams by creating a virtual (layer-3) TUN interface [16] and updating the routing table so that all outgoing traffic, generated by any app on the device, is sent to the TUN interface. The AntMonitor App then routes the datagrams to their target hosts on the Internet (as described below). When a host responds, the response will be routed back to the AntMonitor App, which then sends the response packets to the apps by writing them to TUN.

Traffic Routing. To route IP datagrams generated by the mobile apps and arriving at the TUN interface, the intuitive solution would be to use raw sockets, which unfortunately is not available on non-rooted devices. Therefore, the datagrams have to be sent out using layer-4 (UDP/TCP) sockets, which can be done in one of the following two ways:

1. Client-Server Routing: This follows the design of a typical VPN service: all traffic is routed through a VPN server [14, 13, 15]. The main advantage is the simplicity of implementation: the routing is done seamlessly by the operating system at the server with IP forwarding enabled. However, in a crowdsourcing system with a large number of users, sending all traffic through a VPN server faces scalability challenges. Furthermore, users may not want their traffic to be redirected. Therefore, we used an alternative routing approach that can work entirely on the mobile device, without the need of a VPN server, as described next.

2. Mobile-Only Routing: Routing IP datagrams to target hosts through layer-4 sockets requires a *translation between layer-3 datagrams and layer-4 packets*. For outgoing traffic, data of the IP datagrams has to be extracted and sent directly to the target hosts through UDP/TCP sockets. When a target host responds, its response data is read from the UDP/TCP sockets and must be wrapped in IP datagrams, which are then written to the TUN interface. The Mobile-Only design removes the dependency on a VPN server, thus making AntMonitor self-contained and easy to scale. Furthermore, this design enhances user privacy as all data can now stay on the mobile device and is not routed through a middlebox. Only when the user opts in, select packets are temporarily logged on the AntMonitor App, and then uploaded to the remote LogServer, as described in the next section.

3.3 System Implementation

The AntMonitor system architecture is shown in Fig. 1: it consists of the AntMonitor App and a LogServer.

3.3.1 *On the Mobile:* AntMonitor App

Our current prototype is an Android app[2]. In addition to traffic interception and routing, the app contains a Graphical User Interface (GUI) and modules for logging (Log) and real-time and/or offline Analysis of packets (see Fig. 1).

Fig. 2 shows screenshots of AntMonitor App's **GUI**. It allows the user to turn the VPN service on and off, to

[2]It relies on the VPN service which is available on 95%, *i.e.* billions, of Android devices today. The VPN API has also just been released on iOS version 9.0+ in Sep. 2015; thus, our approach can be implemented on iOS as well.

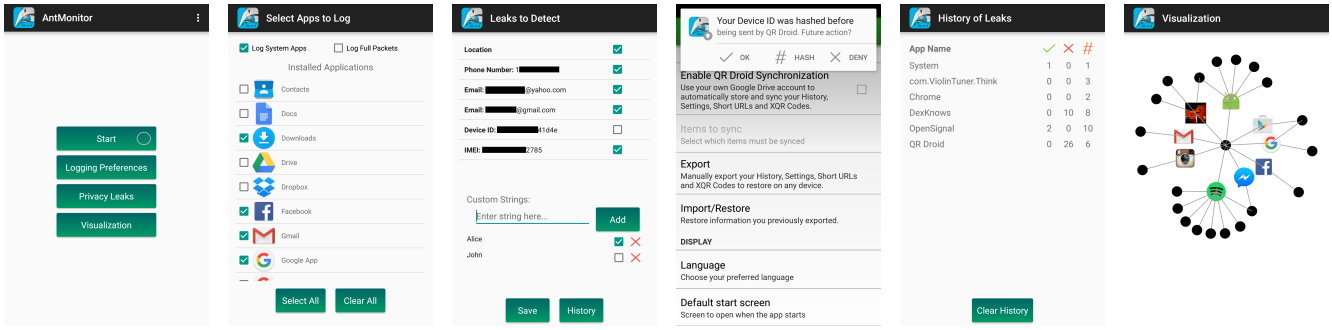

(a) Home Screen (b) Apps to Monitor (c) Strings to Inspect (d) Privacy Leak Alert (e) History of Leaks (f) Visualization

Figure 2: Screenshots of the AntMonitor prototype. (a)(b) apply to all uses of the app. (c-f) are specific to Privacy Leaks. For example, (f) shows which apps leak data to which destinations (typically trackers and ad servers).

select which applications should be monitored and logged; advanced users can also choose to contribute full packets or headers only. As an added benefit, the AntMonitor App shows the results of the applications discussed in Sec. 5. For example, Fig. 2(f) visualizes where each app sends data (*e.g.* to ad servers) by displaying a graph of real-time connections.

The **Forwarder** manages the TUN interface and routes network traffic. It consists of two main components: UDP and TCP Forwarder, both depicted in Fig. 1.

The *UDP Forwarder* is simpler, since UDP connections are stateless. When an app sends out an IP datagram containing a UDP packet, the UDP Forwarder records the mapping of the source and destination tuples (a tuple consists of an IP address and a port number), to be used later for reverse lookup. The Forwarder then extracts the data of the UDP packet and sends the data to the remote host through a *protected* UDP socket. When a response is read from the UDP socket, the Forwarder creates a new IP datagram, changes the destination tuple to the corresponding source tuple in the recorded mapping, and writes the datagram to TUN.

The *TCP Forwarder* works like a proxy server. For each TCP connection made by an app on the device, a TCP Forwarder instance is created. This instance maintains the TCP connection with the app by responding to IP datagrams read from the TUN interface with appropriately constructed IP datagrams. This entails following the states of a TCP connection (LISTEN, SYN_RECEIVED, ESTABLISHED, *etc.*) on both sides (app and TCP Forwarder) and careful construction of TCP packets with appropriate flags (SYN, ACK, RST, *etc.*), options, and sequence and acknowledgment numbers. At the same time, the TCP Forwarder creates an external TCP connection to the intended remote host through a *protected* socket to forward the data that the app sent to the server and the response data from the server to the app.

The **Analysis Module** can do both offline and online analysis on intercepted packets. The online capability allows it to take action on live traffic, *e.g.*, preventing private information from leaking. Since the analyses are done on the mobile, private information is never leaked out of the device, setting AntMonitor apart from systems like Meddle [13], that perform leakage analysis at a VPN server. In order to inspect encrypted traffic, we implement a TLS proxy: see Sec. 5.1.

The **Log Module** writes packets to log files on the device, subject to the user preferences of what apps and what information (*e.g.* entire packets or just packet headers) to log. This module can add rich contextual information to the captured packets by using the PCAP Next Generation format [48]; *e.g.* we currently store application names and network statistics (discussed in detail in Sec. 5.2) alongside the raw packets. The mapping to app names is done by looking up the packets' source and destination IPs and port numbers in the list of active connections available in /proc/net, which provides the UIDs of apps responsible for each connection. Given a UID, we get the corresponding package name using Android APIs. Finally, the Log Module periodically uploads the log files to LogServer over HTTPS, when the device is charging and is connected to Wi-Fi or upon user's request.

3.3.2 *Data Collection Server:* LogServer

The **Log Manager** supports uploading of files using multipart content-type HTTPS. For each uploaded file, it checks if the file is in proper PCAPNG format. If so, for each mobile, the manager stores all of its files in a separate folder.

The **Analysis Module** extracts features from the log files and inserts them into a MySQL database to support various types of analyses. Compared to the Analysis Module of the AntMonitor App, this module (on the LogServer) has access to crowdsourced data from a large number of devices, and can perform global large-scale analyses. For instance, it could detect global threats and outbreaks of malicious traffic.

3.4 Novelty of Design

The key novelty of our design is that we are able to use the minimum number of resources to achieve highest performance. Specifically, we use the minimum number of: (i) threads (two) for network routing, and (ii) sockets.

Thread Allocation. We have fully utilized Java New I/O (NIO) with non-blocking sockets for the implementation of the Forwarder. In particular, Forwarder is implemented as a high-performance (proxy) server, that is capable of serving

hundreds of TCP connections (made by the apps) at once, while using only two threads: one thread is for reading IP datagrams from the TUN and another thread is for actual network I/O using the Java NIO Selector and for writing to TUN. Minimizing the number of threads used is critical on a resource constrained mobile platform so that AntMonitor (which runs in the background) does not interfere with other apps that run in the foreground. For comparison, Privacy Guard creates one thread per TCP connection, which rapidly exhausts the system resources even in a benign scenario, *e.g.*, opening the CNN.com page could create about 50 TCP connections, which results in low performance (see Sec. 4).

Socket Allocation. Since the Forwarder needs to create sockets to forward data and the Android system imposes a limit of 1024 open file descriptors per user process, sockets must be carefully managed. To this end, we minimize the number of sockets used by the Forwarder by (i) multiplexing the use of UDP sockets: we use a single UDP socket for all UDP connections, and (ii) carefully managing the life cycle of a TCP socket to reclaim it as soon as the server or the client closes the connection. For comparison, PrivacyGuard uses 1 socket per UDP connection and 2 sockets per TCP connection; Haystack uses 1 socket per UDP connection and 1 socket per TCP connection.

3.5 Performance Optimization

Since the AntMonitor App processes raw IP datagrams in user-space, it is highly non-trivial to achieve good performance. We investigated the performance bottlenecks of our approach specifically and VPN approaches in general. These bottlenecks are depicted in Fig. 3. We then address these bottlenecks through a combination of techniques, from typical optimization techniques (including implementing custom native C libraries and deploying high-performance network I/O patterns) to highly customized techniques that we specifically devised for the VPN-based architecture. We also provide a detailed comparison of our design to PrivacyGuard [10] (whose source code is publicly available); in contrast, Haystack's source code is unavailable, therefore we qualitatively compare a subset of techniques that we could infer from Haystack's description [11] and our observations.

Traffic Routing (Points 1, 2, and 3). The techniques that we adopted are as follows: (i) we explicitly manage and utilize a Direct ByteBuffer for I/O operations with the TUN interface and the sockets, (ii) we store packet data in byte arrays, and (iii) we minimize the number of copy operations and any operations that traverse through the data byte-by-byte. These techniques are based on the following observations: Direct ByteBuffer gives the best I/O performance because it eliminates copy operations when I/O is performed in native code. Plus, Direct ByteBuffer on Android is actually backed by an array (which is not typically the case on a general Linux platform); therefore, it creates synergy with byte arrays: making a copy of the buffer to a byte array (for manipulation) can be done efficiently by performing a memory block copy as opposed to iterating through the buffer byte-by-byte. (Memory copy is also used whenever a copy of the data is needed, *e.g.*, for IP datagram construction.) Finally, because the allocation of a Direct ByteBuffer is an expensive operation, we carefully manage its life cycle: for an I/O operation, *i.e.*, read from TUN, we reuse the buffer for every operation instead of allocating a new one.

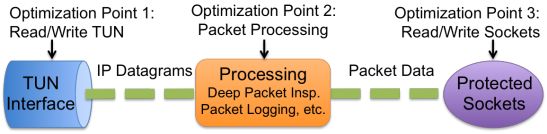

Figure 3: Performance Optimization.

TUN Read/Write (Point 1). The Android API does not provide a way to *poll* the TUN interface for available data. The official Android tutorial [49], as well as PrivacyGuard and Haystack [10, 11], employ periodic sleeping (*e.g.*, 100 ms) between read attempts. This results in wasted CPU cycles if sleeping time is small, or in slow read speed if the sleeping time is large, as the data may be available more frequently than the sleep time. To address this issue, we implemented a *native C library* that performs the native *poll()* to read data to a Direct ByteBuffer (which is then available in the Java code without extra copies).

It is also important to read from (and write to) the TUN interface in large blocks to avoid the high overhead of crossing the Java-Native boundary and of the system calls (read() and write()). In early implementations, we observed that IP datagrams read from the TUN interface have a maximum size of 576 B (which is the minimal IPv4 datagram size). This results in maximum read speed of about 25 Mbps on a Nexus 6 for a TCP connection, thus limiting the upload speed. We were able to increase the datagram size by (i) increasing the MTU of the TUN interface to a large value, *e.g.*, 16 KB and (ii) including an appropriate Maximum Segment Size (MSS) in the TCP Options field of SYN-ACKs sent by TCP Forwarder when responding to apps' SYN datagrams. These changes help to ensure that an app can acquire a high MTU when performing Path MTU Discovery, so that each read from TUN results in a large IP datagram. This results in the maximum read speed, *i.e.*, more than 80 Mbps on our Nexus 6. Similarly, it is important to write to TUN in large blocks: we construct large IP datagrams to write to TUN. We have experimented with other large block values (e.g., 8K, 32K) and found that 16 KB achieved the highest throughput on Nexus 6.

Socket Read/Write (Point 3). Similar to when interacting with the TUN interface, in order to achieve high throughput, it is important to read from (and write to) TCP sockets in large blocks. In particular, we match the size of the buffer used for socket read (*e.g.*, 16 KB minus 40 B for TCP and IP headers) to the size of the buffer used for TUN write (*e.g.*, 16 KB). Similarly, we also matched the size of the buffer used for socket write to that of the buffer used for TUN read. Thus, sending a large IP datagram read from TUN might in-

volve sending several smaller IP datagrams (from the TCP network socket), *i.e.*, fragmentation; however, this is done efficiently in the kernel space. For comparison, PrivacyGuard does not implement this matching.

Packet-App Mapping (Point 2). Android keeps active network connections in four separate files in the /proc/net directory: one each for UDP, TCP via IPv4 and IPv6. Because parsing these files is an expensive I/O operation, we implemented the reading and parsing in a native C library. Furthermore, to minimize the number of times we have to read and parse them, we store the mapping of app names to source/destination IPs and port numbers in a Hash Map. When the Log Module receives a new packet, it first checks the Map for the given IP and port number pair. If the mapping does not exist, the Log Module re-parses the /proc files and updates the Map. Privacy Guard and Haystack also provide packet-app mapping, but their implementation is in Java, and Privacy Guard does not provide logging.

DPI: Deep Packet Inspection (Point 2). Although inspecting every packet is costly, we leverage the Aho-Corasick algorithm [50] written in native C to perform real-time detection without significantly impacting throughput and resource usage (see Sec. 4.1.3). However, this alone is not enough: we must also minimize the number of copies of each packet. Although the algorithm generally operates on Strings, AntMonitor App uses Direct ByteBuffers for routing, and creating a String out of a ByteBuffer object costs us one extra copy. Moreover, Java Strings use UTF-16 encoding and JNI Strings are in Modified UTF-8 format. Therefore, any String passed from Java to native C requires another copy while converting from UTF-16 to UTF-8 [51]. To avoid two extra copies, we pass the Direct ByteBuffer object and let the Aho-Corasick algorithm interpret the bytes in memory as characters. This enables us to perform an order of magnitude faster than Java-based approaches (Sec. 4.3), *i.e.*, compared to Privacy Guard's Java-based string matching and Haystack's Java-based Aho-Corasick implementation.

4. PERFORMANCE EVALUATION

Tool. In order to evaluate AntMonitor, we built a custom app – AntEvaluator. It transfers files and computes a number of performance metrics, including network throughput, CPU and memory usage, and power consumption. It helps us tightly control the setup and compute metrics that are not available using off-the-shelf tools, such as *Speedtest*.

Scenarios. We use AntEvaluator in two types of experiments. In Sec. 4.1, *Stress Test* performs downloads and uploads of large files so that AntMonitor has to continuously process packets. In Sec. 4.2, *Idle Test* considers an idling mobile device so that AntMonitor handles very few packets. In between the two extremes, we have also considered a *Typical Day Test*, which simulates user interaction with apps; however, it is omitted due to lack of space.

Baselines. We report the performance of AntMonitor and compare it to state-of-the-art baselines from Sec. 2.1:

- Raw Device: no VPN service running on the device; this is the ideal performance limit to compare against.
- State-of-the-art mobile-only approaches: Privacy Guard [10] v1.0 and Haystack [11] v1.0.0.8. (We omit the testing of tPacketCapture [17] since it was shown to have very poor performance in [15].)
- Client-server VPN approaches: industrial grade StrongSwan VPN client v1.5.0 with server v5.2.1, and an AntMonitor Client-Server implementation based on [15].[3] The VPN servers used by each app were hosted on the same machine.

Setup. All experiments were performed on a Nexus 6, with Android 5.0, a Quad-Core 2.7 Ghz CPU, 3 GB RAM, and 3220 mAh battery. Nexus 6 has a built-in hardware sensor, Maxim MAX17050, that allows us to measure battery consumption accurately. Throughout the experiments, the device was unplugged from power, the screen remained on, and the battery was above 30%. To minimize background traffic, we performed all experiments during late night hours in our lab to avoid interference, we did not sign into Google on the device, and we kept only pre-installed apps and the apps being tested. Unless stated otherwise, the apps being tested had TLS interception disabled and the AntMonitorApp was logging full packets of all applications and inspecting all outgoing packets. VPN servers ran on a Linux machine with 48-Core 800 Mhz CPU, 512 GB RAM, 1 Gbit Internet; the Wi-Fi network was 2.4Ghz 802.11ac. Each test case was repeated 10 times and we report the average.

4.1 Stress Test

4.1.1 Large File Over a Single Flow

Setup. For this set of experiments, we use AntEvaluator to perform downloads and uploads of a 500 MB file over a single TCP connection. In the background, AntEvaluator periodically measures the following metrics:

A. Network Throughput: AntEvaluator reports the number of bytes transferred after the first 10 sec (to allow the TCP connection to reach its top speed) and the transfer duration. We use these numbers to calculate throughput.

B. Memory Usage: AntEvaluator uses the top command to sample the Resident Set Size (RSS) value.

C. Battery Usage: AntEvaluator uses the APIs available with the hardware power sensor Maxim MAX17050 to compute the energy consumption during each test in mAh [52].

D. CPU Usage: AntEvaluator uses the top command to measure the CPU usage.

At the end of each experiment, AntEvaluator reports the calculated throughput and battery usage, and the average memory and CPU (considering the sum of CPU usage of AntEvaluator and the VPN app) usage.

[3]In the labels of the performance figures, we refer to this baseline for comparison as "AM. Client-Server," to distinguish it from the actual proposed AntMonitor, referred to as "AM. Mobile-Only".

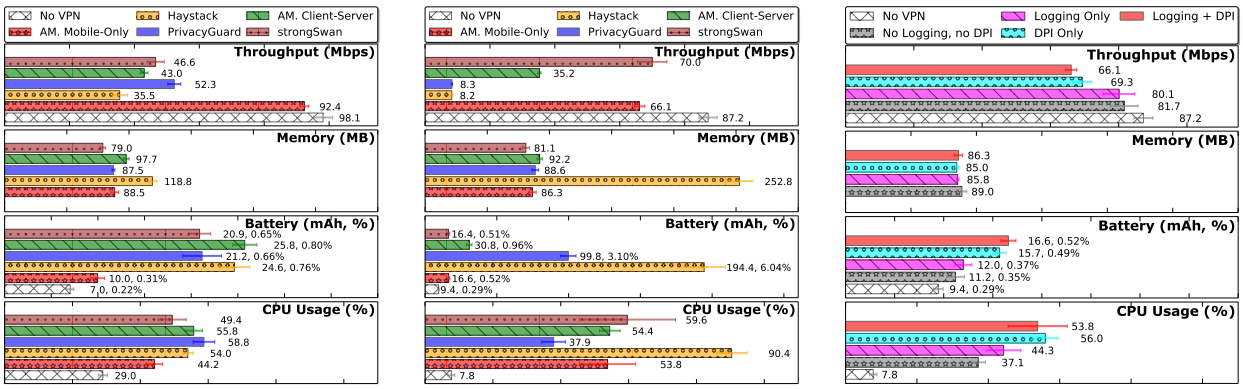

(a) Download (b) Upload (c) 4 Variations of AM. Mobile-Only Upload

Figure 4: Performance of all VPN apps in Stress Test for a 500 MB file on Wi-Fi. ("AM." stands for AntMonitor. "AM. Mobile-Only" stands for AntMonitor proposed in this paper. "AM. Client-Server" is only used as a baseline for comparison.)

Results. Fig. 4(a) shows that the download throughput of AntMonitor significantly outperforms all other approaches. It was able to achieve about *94% of the raw speed*, with throughput *2x more* than StrongSwan & Privacy Guard and *2.6x more* than Haystack. We further note that all VPN apps tested have similar memory, battery, and CPU usage.[4] Fig. 4(b) reports the upload performance. AntMonitor achieves *76% of the raw speed* while performing data logging and DPI. Most significantly, its performance is *8x faster* than both state-of-the-art mobile-only approaches.[5] StrongSwan outperforms AntMonitor as expected since, unlike with incoming packets, AntMonitor performs DPI on each outgoing packet. Nevertheless, Fig. 4(c) shows that AntMonitor has the higher upload speed (and closest to the raw speed) if DPI is disabled. Fig. 4(b) also shows that all VPN apps have similar memory and CPU usage, except for Haystack, which incurs significant overhead. Since the test took longer for the slower approaches, Privacy Guard and Haystack used significantly more battery.

In general, using any VPN service roughly doubles the CPU usage during peak network activity. Although the CPU usage of 38–90% on Wi-Fi seems high, the maximum CPU usage on the quad-core Nexus 6 is 400%. In summary, this set of experiments demonstrates that among all VPN approaches, for both downlink and uplink, AntMonitor has the highest throughput while having similar or lower CPU, memory, and battery consumption.

4.1.2 *Small Files Over Multiple Flows*

Setup. To test the efficiency of AntMonitor's thread and socket allocation, we used AntEvaluator to create 16 threads, each downloading a 50MB file. During the test AntEvaluator calculated the throughput of each flow and reported the average of all flows.

Results. The average speed of a flow (in Mbps) for each test case was the following: Raw Device: 6.82, AntMonitor: 6.57, Privacy Guard: 4.75, StrongSwan: 3.73, Haystack: 3.18, and AntMonitor Client-Server: 3.06. Again, AntMonitor came out on top, achieving 96% of the raw speed.

4.1.3 *Impact of Logging and DPI*

Setup. To assess the overhead caused by Logging Data and Deep Packet Inspection (DPI), we performed the single-flow upload stress test on AntMonitor with all four combinations of Logging on/off and DPI on/off.

Results. First, Fig. 4(c) shows that logging does not have a significant impact on throughput. This is thanks to (i) the optimization of AntMonitor that uses only two threads for network I/O (see Sec. 3.5) and (ii) the fact that the data collection uses two threads for storage I/O. These data logging threads do not significantly impact main network I/O threads on a quad-core Nexus 6 phone. Second, DPI is performed by one of the main network I/O threads and inflicts a 17% slow-down on upload speed. Although 17% is a significant overhead, AntMonitor is still able to reach over 60 Mbps speed, which is more than enough for mobile apps. In addition, DPI causes a 28% and 33% overhead on battery and CPU, respectively. However, the CPU usage still remains 1/8 of the total possible CPU available on the Nexus 6 (of 400%), thus the overhead is acceptable. Finally, without logging and DPI, AntMonitor achieves 94% of the raw speed without VPN.

4.1.4 *Impact of TLS Proxy*

In order to be able to inspect encrypted traffic for privacy leaks, we implemented a TLS proxy, described in Sec. 5.1.

Setup. To evaluate the performance impact of this proxy, we used AntEvaluator to download a 500MB file from a secure server over HTTPS and compared the throughput of AntMonitor to that of the Raw Device.

[4]Although StrongSwan does not perform L3-L4 translation, it performs encryption and decryption, which results in about 5% higher CPU usage than AntMonitor.

[5]The gains provided by the optimizations discussed in Sec. 3.5 have a greater impact on upload speeds because both PrivacyGuard and Haystack favor downstream traffic since the responses from the Internet are read as streams from sockets and the responses from applications are read from TUN packet-by-packet [11].

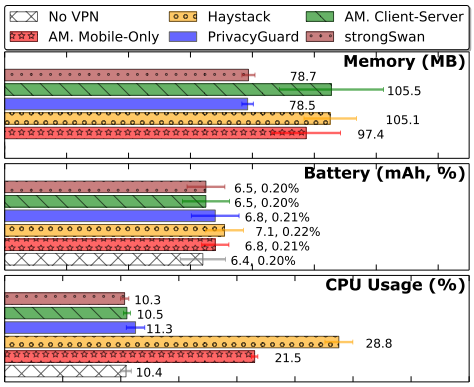

Figure 5: Performance during device idle time.

Results. The average throughput (in Mbps) was 77.2 and 69.1 for the Raw Device and AntMonitor, respectively. As expected, the proxy causes a significant overhead since it uses an extra socket for each connection and performs one extra decryption and encryption operation per packet.

4.2 Idle Test

Setup. For this set of experiments, we kept the phone idle for 2 minutes with only background apps running. We used AntEvaluator to measure the battery and memory consumption of each VPN app. We also measured the aggregate CPU usage across all apps by summing the System and User % CPU Usage provided by the top command.

Results. Fig. 5 shows that all apps tested create very little additional overhead when the device is in idle mode. Among the mobile-only approaches, Haystack and AntMonitor used more CPU than Privacy Guard because both of them have threads to log packets while Privacy Guard does not. Similarly, Logging also results in slightly higher memory usage for Haystack and AntMonitor. (Note that StrongSwan does not log packets either, thus has lower CPU usage.) Finally, the overall memory usage of the VPN apps (~105 MB) is acceptable; many other popular apps, *e.g.* Facebook, use as much as 200 MB of RAM.

4.3 Metrics Computed Outside AntEvaluator

Latency. We measured the latency of each VPN app by averaging over several pings to a nearby server (in the same city). In order of increasing delay, the apps rank as follows: NoVpn: 3 ms, StrongSwan: 4 ms, Haystack: 4 ms, AntMonitor Client-Server: 5 ms, AntMonitor: 7 ms, and Privacy Guard: 83 ms. Compared to client-server approaches, mobile-only approaches cannot forward ICMP packets; thus, we measure latency using TCP packets. The additional delay is due to the time required to create, send, and receive packets through TCP sockets. Compared to Haystack, AntMonitor has a small additional latency as sending and receiving TCP packets involves two threads (one reads/writes the packet from/to TUN and one reads/writes the packet from/to the socket), whereas Haystack might have used a single thread (source code unavailable).

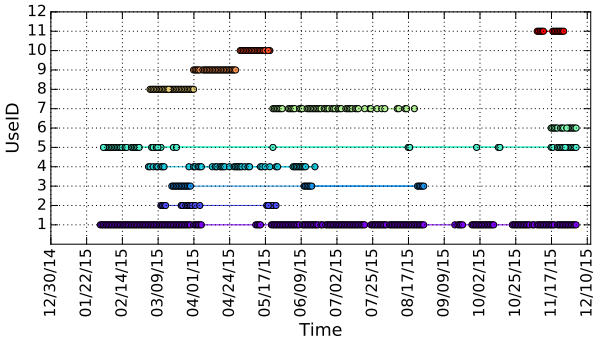

Figure 6: Data logged daily by different users in the study. (We omitted some users, and we have included different devices belonging to the same individual (e.g. 7-11).)

String Parsing. The main heavy operation required in DPI is string parsing. During real traffic conditions, our native C implementation of Aho-Corasick has a maximum run time of 25 ms. When benchmarking as a standalone library (running on the Android main thread alone), our parsing time is <10 ms. For comparison, Haystack reports a 167 ms maximum run time for string parsing with Aho-Corasick.

5. APPLICATIONS

Because AntMonitor intercepts every packet in and out of the device, it is uniquely positioned to serve as a platform for enabling applications that build on top of its passive monitoring capability. We showcase three such applications that may be of interest to operators and/or individual users: (i) privacy leaks detection, (ii) network performance monitoring, and (iii) traffic/user profiling. We report results from a pilot user study at our university campus: Fig. 6 shows the daily activity of 11 (with re-installs) volunteers: these are members of our own research group, who used AntMonitor on their primary phones during the period Feb. 5 – Nov. 30, 2015 and uploaded their data to LogServer. This study was meant only as a proof of concept of the capabilities of AntMonitor and *not* as as representative large scale user study.

5.1 Application I: Privacy Leaks

Mobile devices today have access to personally identifiable information (PII) and they routinely leak it through the network, often to third parties without the user's knowledge. PII includes: (i) mobile phone IDs, such as IMEI (which uniquely identifies a device within a mobile network), and Android Device ID; and (ii) information that can uniquely identify the user (such as phone number, email address, or even credit card) or other information (e.g. location, demographics). Sometimes sending out PII is necessary for the operation of the phone (*e.g.* a device must identify itself to connect to a wireless network) or of the apps (*e.g.* location must be obtained for location-based services). However, the leak may not serve the user (*e.g.* going to advertisers or analytics companies) or may even be malicious. Leaks in plain text can be intercepted by listening third parties, *e.g.* in

App Name	Leak Type	# Flows
VnExpress.net	IMEI, Phone#, Location, Email, DeviceID	33147
Zing Mp3	IMEI, DeviceID	14745
Clean Master	DeviceID	1768
WiFi Maps	IMEI, Location	1582
Relay for reddit	Location	638
System	IMEI, Location, DeviceID	301
Chrome	Location	143
ES File Expl.	DeviceID	96
MyFitnessPal	Location	76
DR Radio	DeviceID	68
Speedtest	Location	48
DexKnows	IMEI,Location	47
Skype	IMEI, DeviceID	42
iWindsurf	IMEI, Location	22
WhatsApp	Phone#	20
...	...	...
All	All	52923

Domain Name	# Flows
eclick.vn.	8760
api.mp3.zing.vn.	7561
ksmobile.com.	620
zaloapp.com.	332
ngoisao.net.	209
api.staircase3.com.	150
openweathermap.org.	47
adtima.vn.	36
ads.adap.tv.	31
apps.ad-x.co.uk.	30
adkmob.com.	27
whatsapp.net.	27
mopub.com.	25
duapps.com.	25
api.dexknows.com.	24
server.radio-fm.us.	15
inmobi.com.	14
...	...
All	18020

Table 1: Flows Leaking PII found in the collected data. (Note: the number of flows at the left (52923 total) is higher than at the right (18020) because we count a flow that leaks multiple PIIs multiple times.)

public WiFi networks. Although modern mobile platforms (Android, iOS, Windows) require that apps obtain permission before requesting access to certain resources, and isolate apps (execution, memory, storage) from each other, this is not sufficient to prevent information from leaking out of the device, *e.g.* due to interaction between apps [53]. Even worse, users are unaware of how their data are used [54].

Privacy Leaks Detection and Prevention Module. We extended the basic AntMonitor functionality with an analysis module that performs real-time DPI. The user can define strings that correspond to private information that should be prevented from leaking; see screenshot in Fig. 2(c). Before sending out a packet, the AntMonitor App inspects it and searches for any of those strings. By default, if the string is found, AntMonitor hashes it (with a random string of the same length, so as not to alter the payload length) before sending the packet out, and asks the user what to do in the future for the given string/app combination, as shown in Fig. 2(d). The user can choose to allow future packets to be sent out unaltered, block them, or keep hashing the sensitive string (so that the application has a good chance to continue working but without the string being leaked). The system remembers the action to take in the future for the same app and "leak," and it will no longer bother the user with notifications. The user may also look at the history of the leaks, shown in Fig. 2(e). AntMonitor can provides both *real-time* detection and *prevention*, on the mobile-device (not at the server) thanks to its efficient implementation, while Haystack and Privacy Guard do not achieve both; and Recon can achieve the same on the server, and could gracefully run on top of AntMonitor as well.

Encrypted Traffic. Since we require plain text in order to perform DPI, and much of the traffic is encrypted, we developed a TLS proxy that uses the SandroProxy library [55], also used by Privacy Guard, to intercept secure connections, decrypt the packets, and then re-encrypt them before sending them to their intended Internet hosts. This method works for most apps, but it cannot intercept traffic

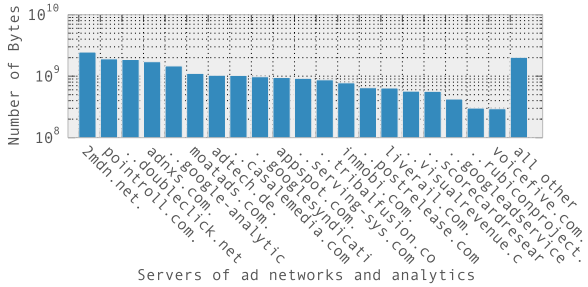

Figure 7: Amount of traffic sent to ad & analytics servers

from highly sensitive apps, such as banking apps, that use certificate pinning. Due to the intrusive nature of intercepting TLS/SSL traffic, we allow users to disable this option.

Privacy Leaks Detected. We analyzed the data collected from our user study, and found a large number of privacy leaks in plain text. Table 1 presents the apps and destination domains with the highest number of flows leaking. The worst offender in the list was the app VnExpress.net that leaked five different types of PII up to 33,145 times towards the domain eclick.vn – an advertising network. The list of leaking apps includes very popular apps with tens of millions of downloads, such as Skype and WhatsApp, and the list of domains includes many mobile ad networks (such as mopub, inmobi, adkmob, adtima). Using the data collected by AntMonitor at the users' mobile devices, we calculated the amount of data that is transmitted from/to ad networks, analytics and mediation services. Fig. 7 shows that the amount of traffic transmitted towards such servers for the top 20 domains is in the order of GBs, consists of several domains, and tens of thousands of flows per domain. Fig. 2(f) visualizes the destinations for one device: it shows which apps leak information to which destinations.

Discussion. The inherent advantage of AntMonitor in this domain is that it runs on the device, which is a more attractive place for users to perform PII scrubbing than a middlebox, since sensitive data does not have to leave the device. Traffic encrypted at the application level can also be intercepted and inspected by the TLS proxy in AntMonitor, then re-encrypted, if the user chooses to do so (by agreeing to install a root certificate). One limitation of our current approach is that PII leakage is detected using string matching, which can be evaded by a more sophisticated attacker, who can leak information across multiple packets or by encoding the text to be leaked. At its current state, our privacy leaks detection is mostly useful for raising awareness of the magnitude, nature and cost of PII leaking, and can be useful in the case of an honest-but-curious adversary: *e.g.*, a legitimate app that leaks information to trackers is likely to stop doing so if users become aware and start uninstalling the app.

5.2 Application II: Performance Measurements

In this section, we demonstrate how to use AntMonitor for monitoring network performance, which can benefit users (*e.g.*, to manage their network access or cost) as well as op-

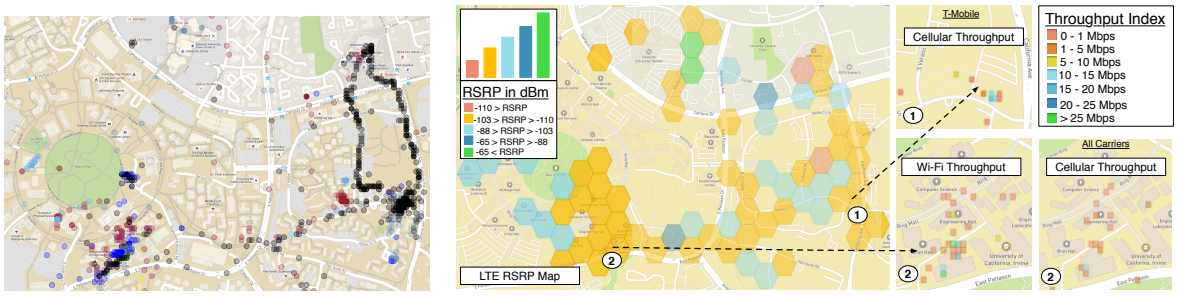

(a) Crowdsourcing (b) T-Mobile LTE Signal Strenth, Wi-Fi & Cellular Throughput Maps

Figure 8: Performance maps from the university campus. (a) Monitor a large area with only a few users. (b) The university has many areas with moderate to poor LTE coverage and many spatial variations. However, low RSRP (location 1) does not necessarily mean low cellular throughput (for the same carrier).

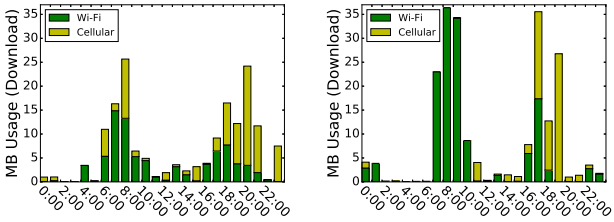

(a) Week Typical Day (b) Weekend Typical Day

Figure 9: **Single user's point of view:** #MB downloaded, averaged over all (a) week or (b) weekend days, for one month, for one user.

Exp #	1	2	3	4	5	6	7
AM: W=1	30.00	32.79	24.49	31.95	34.78	32.31	31.02
AM: W=5	21.24	29.26	22.83	27.01	30.75	26.84	26.14
Speedtest	19.96	28.42	22.39	28.74	31.66	26.98	27.22

Table 2: **Throughput (Download Mbps): Active (using Speedtest) vs Passive (using AntMonitor: AM) measurements**. First, we ran multiple Speedtests, with 5 min gaps, from the same location, and we list the throughput mentioned by Speedtest. Second, we computed the (maximum) average throughput using AntMonitor logs, over a window of 1 & 5 sec.

erators (*e.g.*, to assess and improve their infrastructure [27]). Inherent advantages of AntMonitor in this application domain include the following: (i) passive performance monitoring comes for free (without the need for active probing overhead) since AntMonitor touches every packet transmitted or received; (ii) the ability to correlate network-level metrics (*e.g.*, WiFi and cellular speed) with other information available on the device (*e.g.*, signal strength, type of network, location, time); (iii) the fine-granularity of measurements at the device, app, flow, or destination level.

Performance Module. By design, AntMonitor intercepts every packet and is thus able to passively compute performance indicators of the TCP/IP layer, such as throughput and latency. In addition, it can monitor performance at other layers (*e.g.*, the radio layer) and collect rich contextual information including but not limited to: (i) timestamp; (ii) geolocation in a way that achieves a low energy footprint[6]; (iii) network information (*e.g.*, WiFi or Cellular), radio access technology (RAT), and detailed cellular network information per region (*e.g.*, LTE parameters, frequency bands); (iv) received signal strength (RSS), and (v) throughput and latency measurements per app and overall. For the purposes of this paper, we append RSS, geolocation, RAT, and network connectivity information in PCAPNG files based on network connectivity/location events posted by the OS. The analysis of the collected data is done offline at LogServer.

Example Performance Measurements. In the rest of this section, we report reference signal received power (RSRP) for LTE network, we compute device throughput, and we showcase results for one month (Nov.7- Dec. 7, 2015). Table 2 compares throughput measurements from a state-of-the-art *active* monitoring tool (Speedtest) vs. *passively*. We see that the values they compute are very close, but the passive approach does not incur any measurement overhead. For a fair comparison in this table, we passively monitored the Speedtest packets using AntMonitor. In the wild, throughput computations can be made by counting the number of bytes of actual traffic sent over time. Fig. 9 shows the data (MB) used by one user throughout a typical day (*i.e.*, averaged over all week or weekend days). One can see the breakdown of the traffic into WiFi and cellular, the daily pattern and the difference between week and weekend.[7]

In the future, we plan to use this module of AntMonitor to crowdsource performance measurements, which can provide a comprehensive view of network-wide performance. As an illustrative example, Fig. 8 depicts 5 such performance maps we built for our university campus, reporting signal strength and QoS parameters per cellular carrier and WiFi. Fig. 8(a) depicts the location of 5 users reported during the one month period (Nov. 7-Dec. 7, 2015): only with a few users, we are able to cover a large area of the campus. Network performance varies throughout the day, but thanks to minimal battery overhead, AntMonitor can monitor it continuously and passively, allowing us to collect time-series of network performance. Fig. 8(b) depicts the LTE RSRP for one cellular provider: LTE reception is poor on

[6] AntMonitor listens to location updates by other applications passively and only supplements that with infrequent active requests.

[7] Such statistics are currently reported by mobile devices, but typically at a very coarse granularity (*e.g.*, total amount of data left for this month). In contrast, AntMonitor can report data at fine granularity (e.g. per app, per location, over time, etc.) and can correlate speed with other performance metrics (omitted for lack of space).

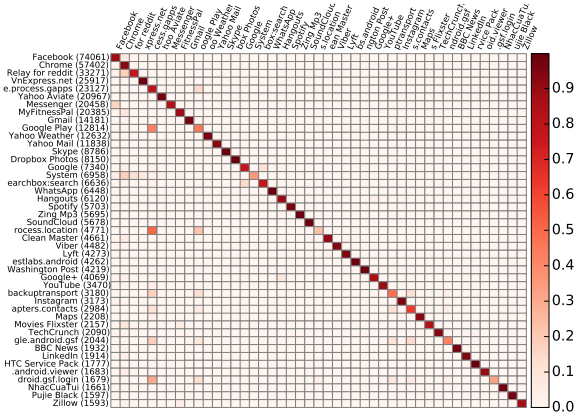

Figure 10: Classification of Flows to Mobile Apps based on TCP/IP features. Normalized confusion matrix for all features. Parentheses show #flows used in training & testing.

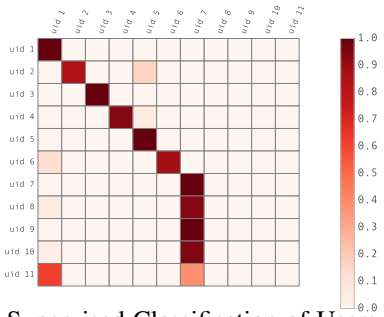

Figure 11: Supervised Classification of Users based on the normalized volume of their mobile apps activity.

many areas and has large spatial variation. Fig. 8(b) also reports the average throughput of WiFi and cellular networks on campus and compares it to LTE RSRP. This information can, for example, be used to inform decisions about switching among, as well as to better provision these networks. Interestingly, in Fig. 8(b), location 1 has low RSRP but high cellular throughput. With the increasing complexity of cellular networks, AntMonitor can help by simultaneously assessing multiple layers of network performance.

5.3 Application III: Learning Profiles

AntMonitor passively monitors all network traffic in and out of the device and can use TCP/IP header features and other contextual information to learn profiles at different levels of granularity, including: per device, app, flow, destination, etc. These can enable anomaly detection, user profiling, market research, traffic differentiation, etc. A key, inherent advantage of AntMonitor, in this application domain, is that it has access to rich contextual information available on the device (*e.g.* location, time, background/foreground apps, application that generated the packets[8]) that can be used to train accurate classification models. We demonstrate two examples: (i) flow classification to mobile apps they belong to, see Sec. 5.4 and (ii) learning user profiles from the apps they use, see Sec. 5.5. These can be useful building blocks for anomaly detection (on the device), traffic differentiation (by the ISP), market research, etc. Training and classification can be performed on the device (Log module in the AntMonitor App) and/or at the LogServer (where data is contributed by multiple devices). In the rest of the section, we report results from the latter.

5.4 Flow classification

[8] For example, although not the first to classify traffic based on packet headers, AntMonitor has an advantage because it operates on the device: for every packet AntMonitor intercepts, it can identify the app it belongs to with 99% accuracy, and append it to the PCAPNG file. This provides ground truth, which is hard to get.

We use packet headers collected in the user study, together with app names that generated the traffic, to train models and classify flows to mobile apps. We used supervised learning to build a multi-class model that classifies network flows to apps. For each flow, we extracted 66 flow features from layer-3 and layer-4 headers (*i.e.* payload and packet size statistics, burstiness, packet inter-arrival times, TCP flags, flow statistics, IP features) on the upstream & downstream directions. We compared different ML models and selected the Random Forest, which performed best. We used a 10-fold cross-validation and kept the same proportions of apps in the testing and training. When we used all features together, the F-1 score was up to 78%. Flow classification for an individual user further increases the classification performance, with the F-1 score ranging between 75%-93%; this is expected, since the number of apps per user is smaller (ranges between 25-70). As a baseline for comparison, random uniform and random proportional classification yield F-1 scores of only 1.5% and 5.8% respectively. Meddle reports a 64.1% precision score in classifying flows for the 92 most popular Android applications by using payload features (Host and User-Agent) [13]. For a dataset of millions of applications, the state-of-the-art approach of AppPrint, that also requires HTTP header data, achieves 81% flow-set coverage with 91% precision. Fig. 10 zooms in the results for the top 45 apps and shows which apps are correctly classified (diagonal entries), while the few errors (non-zero entries off-diagonal) misclassify similar apps to each other (*e.g.*, Facebook to FB Messenger, or Google related apps). The fact that using off-the-shelf learning tools and only features from TCP/IP headers, AntMonitor can classify applications better than state-of-the-art approaches with access to payload, is due to its inherent advantage of having access to accurate ground truth and user behavior at a large scale.

5.5 User Profiling

As a representative example, we asked the following question: can the users in our study group (see Fig. 6) can be distinguished from each based on their app activity? We model each user as the vector of normalized activity volume per app, in one day. One interesting fact in our dataset was that certain users have re-installed AntMonitor during the study

and appear with different user ids. For example, users 7-11 in Fig. 6 are different devices used by the same person over different time periods. First, we use supervised learning in which we include users 1-7 in the training dataset and users 1-11 in testing. Fig. 11 shows the confusion matrix: users 1-7 are correctly classified to themselves, while users 8-10 are classified mostly as user 7, which is also correct. Interestingly, user 11 is classified as user 1, which also makes sense: during that period the 2 users (students in our group) were working on the same deadline and were using their phones for running similar apps for testing and performance evaluation. These results are clearly preliminary but demonstrate AntMonitor's potential for user profiling and anomaly detection, a direction we plan to explore in the future.

6. CONCLUSION

The focus of this paper was on the AntMonitor system for on-device passive mobile network monitoring. Although VPN-based approaches have been used before, AntMonitor's optimized design minimizes the use of resources (*e.g.*, two threads and minimum number of sockets), and significantly outperforms previous state-of-the-art mobile-only approaches, namely Privacy Guard [10] and Haystack [11]: it achieves 2x and 8x faster (down and uplink) speeds and close to the raw no-VPN throughput, while using 2–12x less energy. This significant performance benefit is crucial for the successful adoption of AntMonitor (users are unlikely to install apps that slow down their phones or drain their battery) and for enabling some real-time applications (*e.g.*, privacy leak prevention was not possible before on a mobile-only design).

The applications discussed were meant to showcase the inherent advantage of the AntMonitor system to support them, and they are a standalone topic on their own right. Our pilot deployment at our university campus was also limited to the primary phones of the members of our research group, for alpha testing purposes and as a proof-of-concept; it was not meant as a large scale user study. Despite these limitations, we demonstrated that AntMonitor is a powerful tool for end-users to understand the behavior (at the network, application, and device level) of their device, detect and block privacy leaks, understand where data is transmitted to and correlate patterns; and for operators to correlated network performance measurements at various layers.

In future work, we plan to expand our user base. First, we will make the AntMonitor app available on Google Play. Second, we are developing an SDK that will provide the main functionality depicted in Fig.1 in a compact way to researchers and application developers. We envision that the AntMonitor SDK can provide a crowdsourcing platform for collecting data and enabling privacy, performance and profiling-related research, with a competitive advantage being its optimized design and superior performance. We are working closely with a Mobile Data Analytics company on integrating the AntMonitor SDK in their apps for network performance measurement and market research. We are also collaborating with researchers at ITU, Copenhagen who use AntMonitor in combination with a special wearable to alert users about privacy leaks and investigate HCI issues. We eventually plan to make the SDK available as an open-source library for the research community.

Ongoing systems work includes: providing a modular framework (SDK) for all applications; explicit knobs for configuration (*e.g.*, trading off accuracy for efficiency); and scaling of the LogServer. In terms of applications to explore deeper, we are particularly interested in (i) giving the user control over her personal data (PII leaks) or meta-data (user profile) leaked; and (ii) using AntMonitor to provide QoS at user-space, *e.g.*, prioritizing according to user preferences.

7. REFERENCES

[1] Cisco. The Zettabyte Era—Trends and Analysis. www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/VNI_Hyperconnectivity_WP.html, May 2015.

[2] Q. Xu, J. Erman, A. Gerber, Z. Mao, J. Pang, and S. Venkataraman. Identifying Diverse Usage Behaviors of Smartphone Apps. In *Proc. of ACM IMC*, Berlin, Germany, Nov. 2011.

[3] X. Chen, R. Jin, K. Suh, B. Wang, and W. Wei. Network Performance of Smart Mobile Handhelds in a University Campus WiFi Network. In *Proc. of ACM IMC*, San Diego, California, USA, Nov. 2012.

[4] H. Falaki, D. Lymberopoulos, R. Mahajan, S. Kandula, and D. Estrin. A First Look at Traffic on Smartphones. In *Proc. of the 10th ACM SIGCOMM Conf. on Internet Measurement*, Melbourne, Australia, Nov. 2010.

[5] N. Vallina-Rodriguez, A. Auçinas, M. Almeida, Y. Grunenberger, K. Papagiannaki, and J. Crowcroft. RILAnalyzer: A Comprehensive 3G Monitor on Your Phone. In *Proc. of IMC*, Barcelona, Spain, Oct. 2013.

[6] C. Kreibich, N. Weaver, B. Nechaev, and V. Paxson. Netalyzr: Illuminating the edge network. In *Proc. of the 10th ACM SIGCOMM Internet Measurement Conf. (IMC)*. November 2010.

[7] PhoneLab, University at Buffalo. https://www.phone-lab.org/.

[8] J. Sommers and P. Barford. Cell vs. WiFi: On The Performance of Metro Area Mobile Connections. In *Proc. of ACM IMC*, Boston, USA, Nov. 2012.

[9] X. Wei, L. Gomez, I. Neamtiu, and M. Faloutsos. ProfileDroid: Multi-Layer Profiling of Android Applications. In *Proc. of the 18th ACM annual int. conf. on Mobile computing and networking*, Istanbul, Turkey, Aug. 2012.

[10] Y. Song and U. Hengartner. PrivacyGuard: A VPN-based Platform to Detect Information Leakage on Android Devices. In *Proc. of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices*, Oct. 2015.

[11] A. Razaghpanah, N. Vallina-Rodriguez, S. Sundaresan, C. Kreibich, P. Gill, M. Allman, and V. Paxson. Haystack: In Situ Mobile Traffic Analysis in User Space. *arXiv preprint arXiv:1510.01419*, Oct. 2015.

[12] Networking Group at UC Irvine. The antmonitor project: Crowdsourcing mobile traffic monitoring. http://antmonitor.calit2.uci.edu, 2015.

[13] A. Rao, A. Molavi Kakhki, A. Razaghpanah, A. Tang, S. Wang, J. Sherry, P. Gill, A. Krishnamurthy, A. Legout, A. Mislove, and D. Choffnes. Using the Middle to Meddle with Mobile. Technical report, Northeastern University, Dec. 2013.

[14] strongSwan VPN Client. `https://play.google.com/store/apps/details?id=org.strongswan.android`, 2014.

[15] A. Le, J. Varmarken, S. Langhoff, A. Shuba, M. Gjoka, and A. Markopoulou. Antmonitor: A system for monitoring from mobile devices. In *SIGCOMM Workshop on Crowdsourcing and Crowdsharing of Big Internet Data (C2BID)*, London, UK, Aug. 2015.

[16] Android VpnService. `http://goo.gl/kV7ZZL`, 2014.

[17] tPacketCapture. `www.taosoftware.co.jp/en/android/packetcapture`.

[18] ICSI. Haystack: Understanding the Fate of your Private Data, Project Website and Beta Release. `haystack.mobi`, Oct. 2015.

[19] W. Enck, P. Gilbert, B. Chun, L. P. Cox, J. Jung, P. McDaniel, and A. N. Sheth. Taintdroid: an information-flow tracking system for realtime privacy monitoring on smartphones. *ACM TOCS*, 2014.

[20] A. R. Beresford, A. Rice, N. Skehin, and R. Sohan. Mockdroid: trading privacy for application functionality on smartphones. In *Proc. of HotMobile*, 2011.

[21] P. Hornyack, S. Han, J. Jung, S. Schechter, and D. Wetherall. These aren't the droids you're looking for: retrofitting android to protect data from imperious applications. In *Proc. of CCS*, 2011.

[22] C. Gibler, J. Crussell, J. Erickson, and H. Chen. Androidleaks: Automatically detecting potential privacy leaks in android applications on a large scale. In *Proc. of the Intl. Conf. on Trust and Trustworthy Computing*, 2012.

[23] M. Egele, C. Kruegel, E. Kirda, and G. Vigna. Pios: Detecting privacy leaks in ios applications. In *NDSS*, 2011.

[24] ICSI. Netalyzr for Android. `https://play.google.com/store/apps/details?id=edu.berkeley.icsi.netalyzr.android`.

[25] A. Nikravesh, H. Yao, S. Xu, D. Choffnes, and Z. M. Mao. Mobilyzer: An Open Platform for Controllable Mobile Network Measurements. In *Proc. of the 13th Annual Int. Conf. on Mobile Systems, Applications, and Services (MobiSys)*, Florence, Italy, May 2015.

[26] Alcatel-Lucent. *Motive Wireless Network Guardian*, 2013. `https://www.alcatel-lucent.com/products/wireless-network-guardian`.

[27] C. Cranor, T. Johnson, O. Spataschek, and V. Shkapenyuk. Gigascope: A Stream Database for Network Applications. In *Proc. of the 2003 ACM SIGMOD Int. Conf. on Management of Data*, San Diego, California, USA, June 2003.

[28] A. Gerber, J. Pang, O. Spatscheck, and S. Venkataraman. Speed Testing without Speed Tests: Estimating Achievable Download Speed from Passive Measurements. In *Proc. of the 10th ACM SIGCOMM Internet Measurement Conf. (IMC)*. November 2010.

[29] J. Ren, A. Rao, M. Lindorfer, A. Legout, and D. Choffnes. Recon: Revealing and controlling pii leaks in mobile network traffic. In *Proc. of the 13th Annual Int. Conf. on Mobile Systems, Applications, and Services (MobiSys)*, volume 16, New York, NY, USA, 2016.

[30] Carrier IQ Inc. Mobiles' Diagnostic Analysis. `http://www.carrieriq.com`.

[31] Ookla Speedtest. `https://play.google.com/store/apps/details?id=org.zwanoo.android.speedtest`.

[32] Open Signal Inc. 3G and 4G LTE Cell Coverage Map. `http://www.opensignal.com`.

[33] Sensorly Inc. Unbiased Wireless Network Information. `http://www.sensorly.com`.

[34] Root Metrics Inc. Metro RootScore Reports. `http://www.rootmetrics.com/us`.

[35] Open Signal Inc. US Wi-Fi Report. `https://opensignal.com/reports/2014/us-wifi/`.

[36] Root Metrics Inc. Rootscore Awards and Reports. `http://www.rootmetrics.com/us/rsr/map/2015-1H`.

[37] S. Deng, R. Netravali, A. Sivaraman, and H. Balakrishnan. WiFi, LTE, or Both?: measuring multi-homed wireless internet performance. In *Proc. of the ACM Conf. on Internet Measurement Conf. (IMC)*. 2014.

[38] S. Dai, A. Tongaonkar, X. Wang, A. Nucci, and D. Song. Networkprofiler: Towards automatic fingerprinting of android apps. In *INFOCOM, 2013 Proceedings IEEE*. IEEE, 2013.

[39] S. Miskovic, G. M. Lee, Y. Liao, and M. Baldi. Appprint: Automatic fingerprinting of mobile applications in network traffic. In *Passive and Active Measurement*. Springer, 2015.

[40] H. Yao, G. Ranjan, A. Tongaonkar, Y. Liao, and Z. M. Mao. Samples: Self adaptive mining of persistent lexical snippets for classifying mobile application traffic. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*. ACM, 2015.

[41] T. Karagiannis, K. Papagiannaki, and M. Faloutsos. Blinc: multilevel traffic classification in the dark. In *ACM SIGCOMM Computer Communication Review*, volume 35. ACM, 2005.

[42] T. Karagiannis, K. Papagiannaki, N. Taft, and M. Faloutsos. Profiling the end host. In *Passive and Active Network Measurement*. Springer, 2007.

[43] M. Iliofotou, P. Pappu, M. Faloutsos, M. Mitzenmacher, S. Singh, and G. Varghese. Network monitoring using traffic dispersion graphs (tdgs). In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*. ACM, 2007.

[44] G. Xie, M. Iliofotou, R. Keralapura, M. Faloutsos, and A. Nucci. Subflow: towards practical flow-level traffic classification. In *INFOCOM, 2012 Proceedings IEEE*. IEEE, 2012.

[45] N. V. Verde, G. Ateniese, E. Gabrielli, L. V. Mancini, and A. Spognardi. No nat'd user left behind: Fingerprinting users behind nat from netflow records alone. In *Distributed Computing Systems (ICDCS), 2014 IEEE 34th International Conference on*. IEEE, 2014.

[46] D. Herrmann, C. Gerber, C. Banse, and H. Federrath. Analyzing characteristic host access patterns for re-identification of web user sessions. In *Information Security Technology for Applications*. Springer, 2010.

[47] Android Versions. `developer.android.com/about/dashboards`.

[48] PCAPNG File Format. `http://goo.gl/y89d9U`.

[49] Google, android toyvpn example. `https://android.googlesource.com/platform/development/+/master/samples/ToyVpn/src/com/example/android/toyvpn/ToyVpnService.java`.

[50] Multifast. `http://multifast.sourceforge.net/`.

[51] UTF-8 and UTF-16 Strings. `http://developer.android.com/training/articles/perf-jni.html#UTF_8_and_UTF_16_strings`.

[52] Android BatteryManager API. `http://developer.android.com/reference/android/os/BatteryManager.html`.

[53] A. P. Felt, H. J. Wang, A. Moshchuk, S. Hanna, and E. Chin. Permission re-delegation: Attacks and defenses. In *Proc. of USENIX Security*, 2011.

[54] J. Lin, S. Amini, J. I. Hong, N. Sadeh, J. Lindqvist, and J. Zhang. Expectation and purpose: understanding users' mental models of mobile app privacy through crowdsourcing. In *Proc. of UbiComp*, 2012.

[55] Secure android proxy. `https://code.google.com/archive/p/sandrop/`.